Scalability of IPv6 Transition Technologies for IPv4aaS

draft-lencse-v6ops-transition-scalability

Gábor LENCSE

lencse@sze.hu

Széchenyi István University

IETF 118, v6ops, November 7, 2023.

Aim and Methodology of the Draft

- Aim
  - To provide IPv6 network operators with information about the scalability of different implementations of various IPv4-as-a-Service (IPv4aaS) technologies
- Methodology
  - Carry out RFC 8219 compliant measurements with free software IPv4aaS implementations
  - Publish all the details in research papers
  - Update the Draft with a summary of the most interesting results and add links pointing to the (open access) paper with the details

Progress of the Draft

- Initial version “00” (Presented at IETF 112)
  - Scalability of **`iptables`** stateful NAT44 (served only as an example)
- Versions “01”, “02” (Presented at IETF 113)
  - Scalability of the **Jool** Stateful NAT64 implementation
    - Performance characteristics:
      1. maximum connection establishment rate (new, stateful specific metric)
      2. throughput (classic RFC 2544 / RFC 8219 metric)
      3. connection tear down rate (new, stateful specific metric)
    - Measured against
      - the number of CPU cores: 1, 2
      - the number of concurrent sessions: 1, 2, 3

Progress of the Draft

- Version “03” (not presented)
  - Connection tracking table capacity measurements of **`iptables`**
  - Connection establishment validation of Jool
    - They both were interesting only from measurement methodological point of view
- Version “04” (Presented at IETF 115)
  - Scalability of the **Jool** implementation of 464XLAT and MAP-T
    - Measured with DNS traffic using the dns64perf++ measurement tool
    - The measurements were not RFC 8219 compliant, but gave an important insight
- Version “05” (current version)
  - Added: scalability measurements of OpenBSD PF

Why OpenBSD PF?

- OpenBSD focuses on security

  Only two remote holes in the default install, in a heck of a long time!

- PF (Packet Filter) is the firewall of OpenBSD
  - Supports stateful NAT64 since OpenBSD 5.1
- The tested version was OpenBSD 7.1
  - Having PF as single threaded, thus scalability was not tested against the number of CPU cores

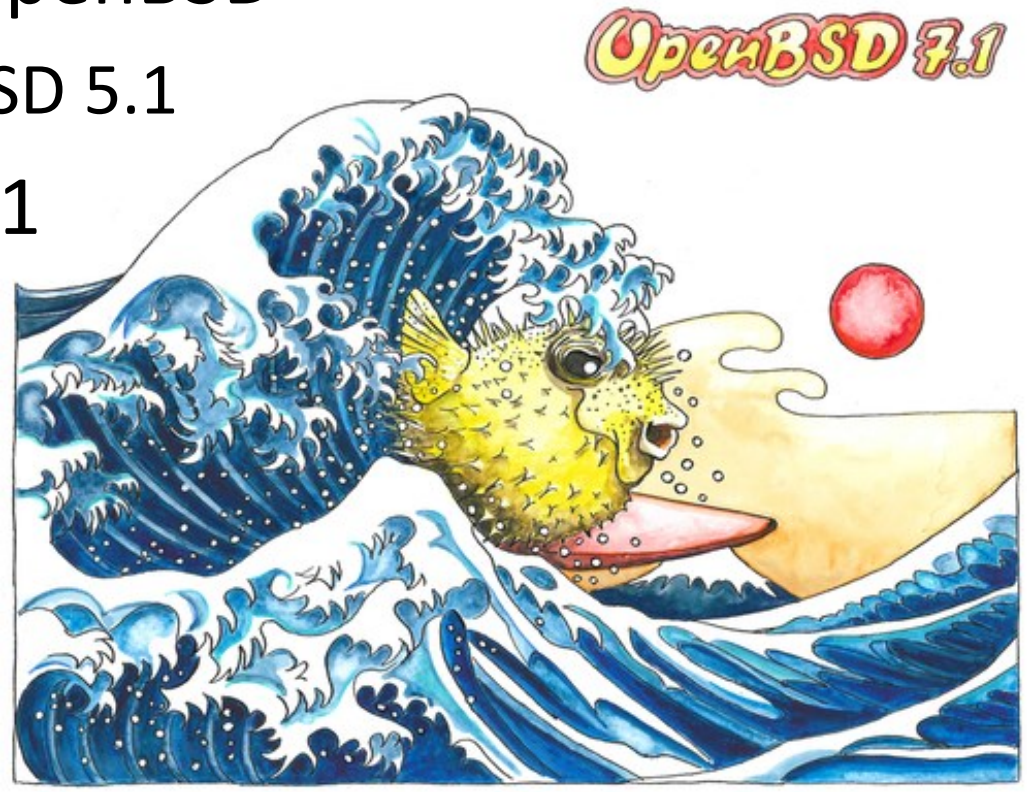

Measurement Environment

- Dell PowerEdge R430 servers
  - Intel Xeon E5-2683v4 CPUs,
  - 384GB 2400MHz DDR4 RAM
  - Intel 10G dual-port X540 NIC
- Direct cable connections
- Tester: Debian 9.13
  - https://github.com/lencsegabor/siitperf
- DUT: OpenBSD 7.1

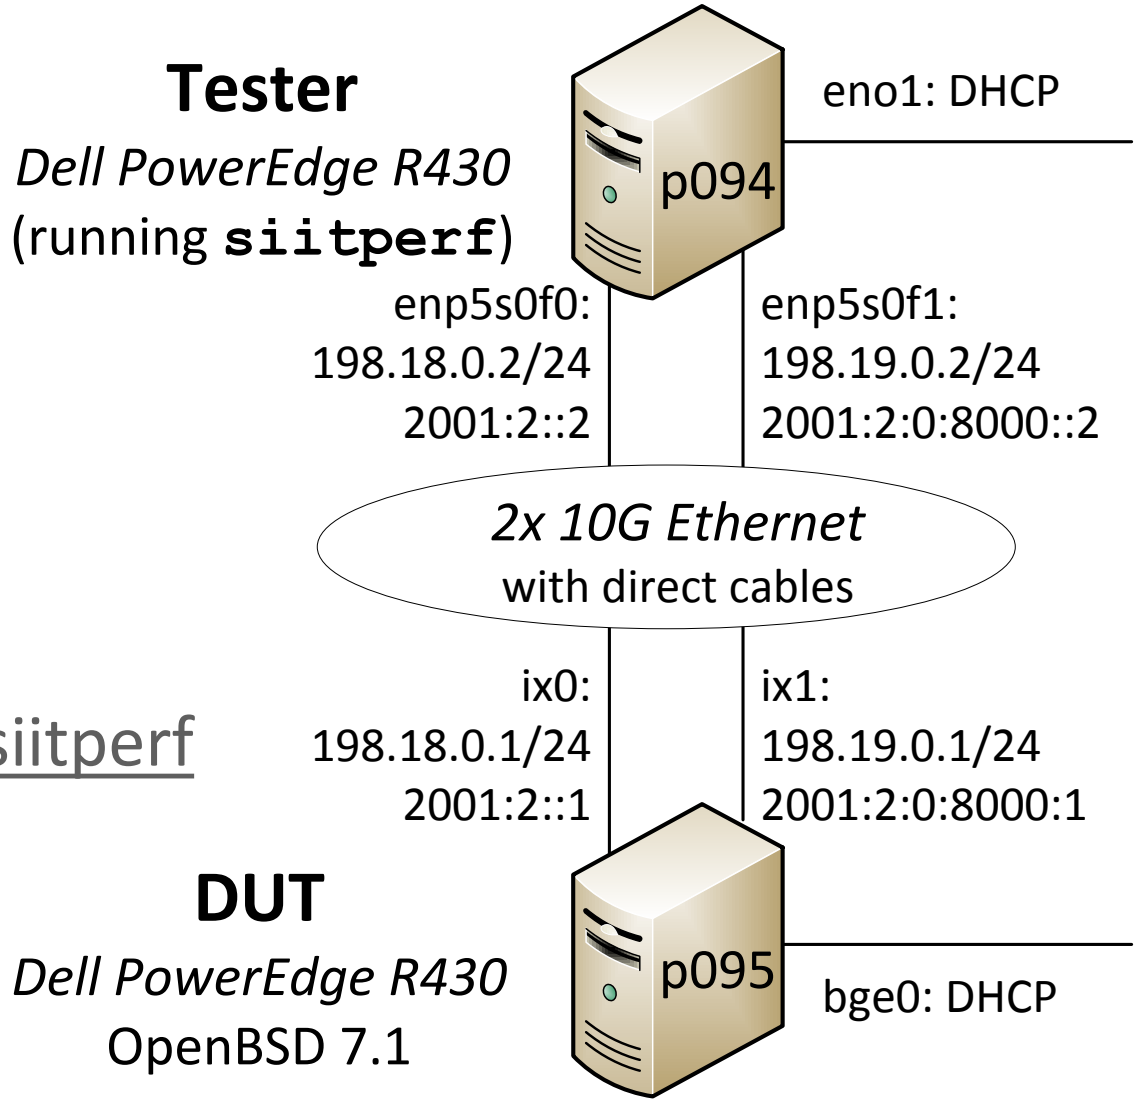

Scalability of OpenBSD PF

- **Maximum Connection Establishment Rate*** of OpenBSD PF as a Function of the *Number of Connections*

Number of connections	400,000	4,000,000	40,000,000
Source port numbers	40,000	40,000	40,000
Destination port numbers	10	100	1,000
Error (cps)	50	40	50
Median (cps)	120,214	85,039	74,022
Minimum (cps)	118,701	84,882	73,680
Maximum (cps)	122,411	85,351	74,266
Median / previous median	-	0.71	0.87

  – The results show an acceptable performance degradation

**This is a new, stateful-specific performance metric*

Scalability of OpenBSD PF

- **Throughput*** of OpenBSD PF as a Function of the *Number of Connections*, Bidirectional Traffic, number of all packets

Number of connections	400,000	4,000,000	40,000,000
Source port numbers	40,000	40,000	40,000
Destination port numbers	10	100	1,000
Error (fps)	200	80	100
Median (fps)	237,304	198,828	173,338
Minimum (fps)	236,912	198,046	172,946
Maximum (fps)	250,584	199,452	174,120
Median / previous median	-	0.84	0.87

  – The results show an acceptable performance degradation

**This is a classic RFC 2544 / RFC 8219 performance metric*

Scalability of OpenBSD PF

- **Connection Tear Down Rate*** of OpenBSD PF as a Function of the *Number of Connections*, 16 CPU Cores (**new, stateful-spec.*)

Number of connections	400,000	4,000,000	40,000,000
Source port numbers	40,000	40,000	40,000
Destination port numbers	10	100	1,000
Filled table del. time med. (s)	1.45	11.56	94.20
Filled table del. time min. (s)	1.36	11.03	91.73
Filled table del. time max. (s)	1.78	13.81	118.52
Empty table del. time med (s)	0.37	0.37	0.37
Empty table del. time min (s)	0.36	0.36	0.36
Empty table del. time max (s)	0.37	0.37	0.37
Connections deletion time (s)	1.08	11.19	93.83
Connection tear down r. (cps)	370,370	357,622	426,303

Are you interested in the details?

- All the details and further results can be found in:

  G. Lencse, K. Shima, K. Cho, "Benchmarking methodology for stateful NAT64 gateways", *Computer Communications*, vol. 210, October 2023, pp. 256-272, DOI: http://doi.org/10.1016/j.comcom.2023.08.009

  (open access paper)

- Acknowledgements

  - This work was supported by the International Exchange Program of the National Institute of Information and Communications Technology (NICT), Japan.
  - The experiments were carried out remotely using the resources of NICT StarBED, 2-12 Asahidai, Nomi-City, Ishikawa 923-1211, Japan.

Questions?

- Do you have any questions or concerns to be addressed?
- Do you consider this draft useful?
- Do you think that this draft should be adopted as a WG item?